arXiv:1605.07369v1 [cs.CV] 24 May 2016

Quickest Moving Object Detection

Dong Lao and Ganesh Sundaramoorthi
King Abdullah University of Science & Technology (KAUST), Saudi Arabia
{dong.lao,ganesh.sundaramoorthi}@kaust.edu.sa

Abstract

We present a general framework and method for simultaneous detection and segmentation of an object in a video that moves (or comes into view of the camera) at some unknown time in the video. The method is an online approach based on motion segmentation, and it operates under dynamic backgrounds caused by a moving camera or moving nuisances. The goal of the method is to detect and segment the object as soon as it moves. Due to stochastic variability in the video and unreliability of the motion signal, several frames are needed to reliably detect the object. The method is designed to detect and segment with *minimum* delay subject to a constraint on the false alarm rate. The method is derived as a problem of *Quickest Change Detection*. Experiments on a dataset show the effectiveness of our method in minimizing detection delay subject to false alarm constraints.

1 Introduction

Segmentation of object(s) from video is a fundamental problem in computer vision. Motion cues play a role in biological visual systems, and they may play an important part in segmentation of objects in both biological and computer vision systems [1]. Thus, there have been many works that segment a video by apparent motion in an attempt to segment relevant objects (e.g., [2, 3, 4, 1, 5]). In many of these methods, it is assumed that the video is observed when the objects of interest are already in motion and the task is to segment such objects. However, that may not be representative of the problem solved by biological systems or that is required in certain computer vision applications such as surveillance. In such cases, the object may be stationary or out of the field of view of the observer or camera when the video starts. Thus, *detection* and segmentation of the object at the time it moves is needed.

This paper addresses the problem of *detection* and *segmentation* of an object in a video that moves, at some unknown time, differently than a dynamic "background", induced from camera motion or slight moving nuisances in the scene, such as leaves in a tree blown by wind. Since we wish to address a problem that is closer to the problem that may be solved by biological systems, an *online* algorithm is desired. We define an online system as a system that receives frames sequentially, one at a time, and must make a decision, that is, declare a detection or wait for more data, at each time instant. Although observing mores frames before declaring a detection may lead to a more accurate detection and segmentation, since more motion may be observed leading to a stronger motion cue and more data mitigates the stochastic effects of background nuisances, this leads to greater delay, which may not be tolerable in certain scenarios, e.g., an antelope detecting an approaching lion, or a surveillance system detecting and tracking an intruder. Thus, our goal is to derive an algorithm with *minimum delay*, defined as the number of frames acquired after the object moves. Of course one can easily achieve zero delay by always declaring detection at frame 1, irrespective of the data. Thus, we require an algorithm that operates under a constraint on *false alarms*, defined as declarations of detection before the object moves or incorrect or inaccurate segmentation after the object moves.

Our specific contributions are as follows: **1.** We derive statistical models of image sequences for the preceding problem so as to formulate the problem in the framework of *Quickest Change Detection*

(QCD) considered in statistics [6, 7]. That problem assumes a stochastic process, which arises from a given distribution before an *unknown* change time and a different distribution post-change. That theory derives online algorithms to detect the change time with *optimal* delay subject to false alarm constraints. **2.** We illustrate and address additional complications in QCD arising from moving object detection, namely, that pre- and post-change distributions depend on hidden variables (motions, occlusions, segmentations), which must be estimated from non-convex optimization problems. **3.** We derive an online algorithm for detection and segmentation of a moving object based on QCD, inheriting *optimality* properties. Further, we introduce an application specific modification of the algorithm to address computational cost of the algorithm derived from QCD. **4**. Finally, we quantitatively evaluate the algorithm on a dataset and compare it to a baseline approach in terms of the desired optimality.

1.1 Related Work

We briefly highlight related work in motion segmentation, and change detection. Motion segmentation relies on computing apparent motion, determined from parametric models (e.g., translations or affine) of flow [8] or dense optical flow [9, 10, 11]. Early works (e.g., [2, 3]) on motion segmentation use parametric models of flow and solve a joint problem in segmentation and flow estimation. To deal with deforming objects such as articulating humans, non-parametric motion models of flow are assumed (e.g., [12]), and solved as a joint problem of segmentation and flow estimation. [4, 5] use a similar approach to causally segment videos frame-by-frame. Those approaches typically operate on 3 frames. To obtain a stronger motion signal, whole videos are processed in batch, rather than online, in [1, 13] by grouping trajectories of points across frames. Although the causal approach [5] is online and uses few frames, the motion signal may not be strong enough to segment the object. On the other hand, [1, 13] may achieve a stronger motion signal at the expense of processing the whole batch. Thus, an algorithm that chooses the number of frames to achieve a strong enough motion signal would be beneficial, and our approach encompasses this. Moreover, existing approaches, while addressing segmentation, do not address *detection* of a moving object at an unknown time in the video, which is the main motivation of our method.

The problem of detecting changes in a video has a large literature in computer vision, too extensive to review here, and thus we refer to [14] for a good survey. That literature mainly addresses detection and segmentation of moving objects by background subtraction [15] - subtraction of a known or adaptively determined background from the current video frame. We are interested in video with moving cameras or dynamic background nuisances, for which the methods in that literature largely do not apply, although there have been advancements. While there are methods that deal with dynamic cameras and detect and segment moving objects by motion (e.g., [16]), they do not address the issue of the tradeoff between detection delay and false alarms, which our approach addresses optimally by using QCD. We are not aware of the use of QCD in the detection and segmentation by motion of video. In the next section, we summarize the main ideas from QCD before framing our problem in that framework.

2 Overview of Quickest Change Detection Theory

We introduce the main theory, *Quickest Change Detection* (QCD) [6, 7], underpinning our algorithm for moving object detection. QCD considers the problem of *sequentially* detecting changes in distribution of a discrete-time stochastic process $\{X_i\}_{i=1}^{\infty}$. It is assumed that X_i is sampled from a known distribution p_0 before an unknown *change time* Γ, and that X_i is sampled from a known distribution p_1 at and after the change time. The literature on QCD derives algorithms for determining the change with fewest observations X_i after the change subject to constraints on false alarms (a false alarm is a declared change by the algorithm before the change time Γ). The main motivation of this theory is that due to the stochastic nature of the problem, algorithms attempting detection of changes sequentially by seeking changes in adjacent samples may lead to false alarms and significant performance gains can be obtained by incorporating all samples before the current time. Of course a strategy that uses more samples would reduce false alarms, but lead to greater delay in detecting the change. QCD addresses this tradeoff between false alarms and number of samples after the change time *optimally*.

QCD is formulated as an optimization problem. Based on the priors available on the change time Γ, the optimization problem can be formulated as a Bayesian or minimax problem. Since we do not have such a prior in our application, we use a minimax formulation that we review from [7] next. A *stopping time* τ with respect to a stochastic process $\{X_i\}_{i=1}^{\infty}$ is a random variable such that the event $\{\tau = n\}$ is in the sigma-algebra generated by $X_1, \ldots, X_n$. Intuitively, τ is a function that may return n if it uses only information determined from $X_1, \ldots, X_n$. An example is $\tau = \inf\{n : \sum_{i=1}^{n} X_i \geq b\}$, i.e., τ is the first time n that the sum of X_i up to time n exceeds a threshold b. Let $\mathbb{P}_n$ and $\mathbb{E}_n$ denote the probability measure and expectation, associated with a change time of n. The *average detection delay* of a stopping time τ in a minimax formulation is defined as

$$\mathsf{ADD}(\tau) = \sup_{n \geq 1} \mathbb{E}_n[\tau - n | \tau \geq n]. \tag{1}$$

The delay is $\tau - n$ given that the change time is n and $\tau \geq n$. ADD defines the worst case average delay over all change times. The *false alarm rate* of a stopping time is defined as $\mathsf{FAR}(\tau) = 1/\mathbb{E}_{\infty}[\tau]$, that is, one over the average stopping time given that there is no change. Note that there are different ways of defining the average detection delay and false alarm rate, but all lead to similar optimal stopping times. The QCD optimization problem is

$$\min_{\tau} \mathsf{ADD}(\tau) \text{ subject to } \mathsf{FAR}(\tau) \leq \alpha, \tag{2}$$

where $\alpha \in [0, 1]$ is the maximum tolerable false alarm rate. It can be shown that the optimal stopping time is given by a maximum likelihood ratio test:

$$\tau_c = \inf\{n : \Lambda_n \geq b\}, \tag{3}$$

where the threshold depends on α, and the *likelihood* Λ_n arises from a test of the null hypothesis $\Gamma \geq n$ against the alternative hypothesis $\Gamma < n$, i.e.,

$$\Lambda_n = \mathbb{P}[\Gamma < n | X_1, \ldots, X_n] / \mathbb{P}[\Gamma \geq n | X_1, \ldots, X_n] = \max_{1 \leq k < n} \prod_{i=k}^{n} p_1(X_i)/p_0(X_i), \tag{4}$$

where the last equality is made under the assumption that X_i are iid before and after the change. In our moving object detection problem that we introduce next, the post-change distribution is only known conditional on a parameter θ, and in this case, the stopping time is $\Lambda_n = \max_{1 \leq k < n} \max_{\theta} \prod_{i=k}^{n} p_\theta(X_i)/p_0(X_i)$, which also optimizes (2).

The sequential algorithm to detect the change is to acquire data $X_i = x_i$ and, at each new acquisition at time n, one computes Λ_n, by solving a maximization problem over possible change times from $k = 1, \ldots, n-1$. At the first n when $\Lambda_n \geq b$, a detection is declared and the maximizer k_n^* is the change time. Solving the maximization problem directly may be expensive. In the case that the post-change distribution is known, previous data does not need to be revisited and Λ_n can be computed recursively. However, the case of the post-change distribution that has an unknown parameter does not have a recursive implementation. We suggest later an approach to reduce the cost of evaluating the likelihood for all $1 \leq k < n$ at each arrival of a new sample at time n.

3 Quickest Moving Object Detection

In this section, we formulate the problem of sequentially detecting and segmenting moving objects from a video quickly as a Quickest Change Detection problem. We first present our statistical models for the image sequence and moving object, then formulate them in the framework of Quickest Change, show the additional complications that arise from our models, and then propose an algorithm.

3.1 Model for the Image Sequence and the Moving Object

We assume that the scene is observed by a possibly moving observer, i.e., the "background" may be moving and at some time Γ an object within the scene begins to move or comes into view of the camera. Although our algorithm will apply to more than one object that may occur at different change times, for simplicity of notation in deriviing our models, we assume the case of one moving object at one change time.

To apply Quickest Change Detection, we specify statistical models pre- and post-change. Let $\Omega \subset \mathbb{R}^2$ be the domain of the images, and let $I_i : \Omega \to \mathbb{R}^k$, $i \geq 1$ be the image sequence. We denote by

$R_i \subset \Omega$ the region of the moving object at time $i \geq \Gamma$. We set $R_i = \emptyset$ for $i < \Gamma$. We denote $O_i^1 \subset R_i$ as the occlusion of the object induced by a change of viewpoint of the camera or a self-occlusion. Similarly, $O_i^0 \subset \Omega \backslash R_i$ is the occlusion of the background induced by the previous conditions or from occlusion induced by the moving object.

The *background displacement* between two adjacent frames in the image sequence is $v_i^0 : \Omega \backslash R_i \backslash O_i^0 \to \mathbb{R}^2$, and the *object displacement* is $v_i^1 : R_i \backslash O_i^1 \to \mathbb{R}^2$. Although such displacements are not defined in the occluded parts of the domain, they will be smoothly extended into the occlusion. We denote the *warp* between frames i and $i+1$ as $w_{i,i+1}^0 : \Omega \backslash R_i \backslash O_i^0 \to \Omega \backslash R_{i+1}$ and $w_{i,i+1}^1 : R_i \backslash O_i^1 \to R_{i+1}$, defined by $w_{i,i+1}^j(x) = x + v_i^j(x)$, which are diffeomorphisms onto their ranges, which arise from camera motion of a non-planar scene or from deforming objects. The warp between time 1 and time i will be denoted by w_i^j and can be determined from the recursion (we assume $v_i^1 = 0$, $i < \Gamma$)

$$w_{i+1}^j(x) = w_i^j(x) + v_i^j(w_i^j(x)), \ i > 1 \quad \text{with} \quad w_1^j(x) = x. \tag{5}$$

The model for the evolution of the region across time after the change time is

$$R_{i+1} = w_{i,i+1}^1(R_i) \cup D_{i+1} \text{ for } i \geq \Gamma \quad \text{and} \quad R_\Gamma = R, \tag{6}$$

where $D_{i+1} \subset \Omega$ is the part of the object that is disoccluded (comes into view of the camera) at time $i+1$, and R is the region of the object at the change time Γ. Therefore, the region of the object is a smooth warping of an initial region R (up to disocclusions), and therefore also smoothly varies in time.

Assuming approximate Lambertian reflectance of the scene, we may relate successive images before the change as

$$I_{i+1}(w_{i,i+1}^0(x)) = I_i(x) + \eta_i(x), \quad x \in \Omega \backslash O_i^0, \ i < \Gamma, \tag{7}$$

and after the change time as

$$\begin{cases} I_{i+1}(w_{i,i+1}^0(x)) = I_i(x) + \eta_i(x) & x \in \Omega \backslash R_i \backslash O_i^0 \\ I_{i+1}(w_{i,i+1}^1(x)) = I_i(x) + \eta_i(x) & x \in R_i \backslash O_i^1 \end{cases}, \ i \geq \Gamma, \tag{8}$$

where $\eta_i(x)$ is a Gaussian iid noise process in both i and x, which is used to model deviations from the Lambertian assumption. We assume that $\eta_i(x) \sim \mathcal{N}(0, \sigma_{\eta,1})$ for $x \in R_i \backslash O_i^1$ and $\eta_i(x) \sim \mathcal{N}(0, \sigma_{\eta,0})$ for $x \in \Omega \backslash R_i \backslash O_i^0$.

3.2 Algorithm for Detection and Segmentation

We now apply the theory in Section 2 to the models of the previous sub-section to derive an algorithm that detects and segments the object R_i at the detection time. We must compute the likelihood ratio, Λ_n, which requires the computation of the pre- and post-change distributions. Our model is complicated by hidden variables, R_i, O_i^j, v_i^j, which must be marginalized out to determine the probabilities. We may treat these distributions that depend on hidden variables as a post-change distribution with an unknown parameter θ, described in Section 2: we maximize over the hidden variables.

Let $\mathbf{I}_{1:n}$ denote all the images $I_1, \ldots, I_n$, and similarly define $\mathbf{v}_{1:n}^j$ and $\mathbf{R}_{1:n}$. Note that conditional on v_i^j and R_i, the pairs of images I_i, I_{i+1} are independent for all i. Using this, one can show that pre-change,

$$p_0(\mathbf{I}_{i:i+1} | v_i^0) \propto \exp\left\{ -\frac{1}{2\sigma_{\eta,0}^2} \int_\Omega \rho(I_{i+1}(w_{i,i+1}^0(x)) - I_i(x)) \, \mathrm{d}x \right\}, \tag{9}$$

where $\rho(x) = \min\{x^2, \beta\}$ ($\beta > 0$) is a truncated quadratic, which is a robust norm that eliminates the explicit estimation of the occlusion [10]. Similarly, the post-change conditional distribution is

$$p_1(\mathbf{I}_{i:i+1} | v_i^0, v_i^1, R_i) \propto \exp\left\{ -\sum_{j=0}^{1} \frac{1}{2\sigma_{\eta,j}^2} \int_{R_i^j} \rho(I_{i+1}(w_{i,i+1}^j(x)) - I_i(x)) \, \mathrm{d}x \right\}, \tag{10}$$

where $R_i^1 = R_i$ and $R_i^0 = \Omega \backslash R_i$. By maximizing over the conditioned variables, one can use the right-most equality in (4) to deduce

$$-\log \Lambda_n = \min_{1\le k<n} \min_{\mathbf{v}^0_{k:n}} \sum_{i=k}^{n} \int_\Omega \text{Res}_i^0(x)\,\mathrm{d}x - \min_{\mathbf{R}_{k:n}} \sum_{j=0}^{1} \min_{\mathbf{v}^j_{k:n}} \sum_{i=k}^{n} \int_{R_i^j} \text{Res}_i^j(x)\,\mathrm{d}x - \log\left[p(R_i)\right], \quad (11)$$

where $\text{Res}_i^j(x) = \frac{1}{2\sigma^2_{\eta,j}} \rho(I_{i+1}(w^j_{i,i+1}(x)) - I_i(x))$.

Warp estimation: To evaluate Λ_n, we must solve an optimization problem in the change time k, and the hidden variables, which we describe next. Given the other variables, we discuss the optimization in $\mathbf{v}^j_{k:n}$. Note that no explicit prior on the warps are assumed in the marginalization of the conditional probabilities, and thus no regularization of the warps appear in the optimization for Λ_n. Instead, we leverage on the Sobolev framework [5], to impose regularity in a natural coarse-to-fine framework allowing warps to be arbitrary diffeomorphisms by constructing warps that are a time integration of smoothly varying vector fields, which at each instant belongs to a Sobolev space. This avoids the under/over smoothing problem in global regularization typically used in optical flow. Solving for v_i^j is done as $v_i^j = \arg\min_v \int_{R_i^j} \rho(I_{i+1}(x + v(x)) - I_i(x))\,\mathrm{d}x$, as other terms are independent of v_i^j.

Region estimation: Given estimates of the warps, we optimization for $\mathbf{R}_{k:n}$, which involves the second summation in (11). Note that R_i are coupled through (6). In the case of no disocclusions, optimizing in $\mathbf{R}_{k:n}$ can be replaced by optimization in any one region R_m with $k \le m \le n$ subject to the constraint that the other regions are warps of R_m. This gives the minimizer of Λ_n as the minimizer of

$$E(R_m) = \sum_{j=0}^{1} \int_{R_m^j} f_i^j(x)\,\mathrm{d}x - \log\left[p(R_m)\right], \quad f^j(x) = \sum_{i=k}^{n} \rho(I_{i+1}(w^j_{m,i+1}(x)) - I_i(w^j_{m,i}(x))), \quad (12)$$

where $w^j_{m,i}$ denotes the warp of R_m^j to R_i^j, determined by (5) using given estimates of $\mathbf{v}^j_{k:n}$. Due to the aperture problem and occlusions, motion residuals in these cases are unreliable to estimate R_m, and we rely on image statistics as in [5]. Thus, we modify the energy as

$$E_{seg}(R_m) = \sum_{j=0}^{1} \int_{R_m^j} (1 - \text{maf}(x)) f^j(x) - \text{maf}(x) \log p_{R_m^j}(I_m(x))\,\mathrm{d}x - \log\left[p(R_m)\right], \quad (13)$$

where $\text{maf} : \Omega \to [0,1]$ is the *motion ambiguity function*, and $p_{R_m^j}$ are local color histograms of I_m within regions. The motion ambiguity function is 1 if pixel $x \in R_m$ is occluded in *all* frames $i = k, \ldots, n, i \neq m$ (all high residuals exceed β), or if x is in a textureless sub-region of I_m. In these cases, motion cues are unreliable, and thus we rely on color histograms to group x. We choose the prior $p(R_m)$ to induce smoothness, as is typical in segmentation problems. The energy (13) now fits into a form considered in [5]. Thus, we use the optimization specified there, which uses gradient descent due to non-convexity. Even though we assumed no disocclusions, optimization of (13) implicitly computes disocclusion as part of the grouping procedure. In particular, disocclusions (parts of the object that come into view) at time m are assumed to be parts of the image moving similar to R_m (or similar color intensity in the case of motion unreliability). Note that the although the optimization problem is in R_m, each of R_i for $i \in \{k, \ldots, n\} \backslash \{m\}$ can be determined as a propagation of R_m through the sequence determined by warps $w^j_{i,i+1}$, which when done through [5], includes disocclusion in the R_i.

Joint Region and Warp Estimation: For a given change time k and time n, we state our algorithm to estimate the detected region $R_{k,n}$ in frame n, as well as the inner optimization in (11), which we denote $\log \Lambda_{k,n}$. Although we only seek $R_{k,n}$ and $\Lambda_{k,n}$, we must estimate the nuisance variables $\mathbf{v}^j_{k:n}$ and $\mathbf{R}_{k:n}$. The algorithm is stated in Algorithm 1. The gradient descent of E_{seg} requires an initialization for R_m. We initialize it with a clustering of the cumulative optical flow from frame m to n (and frame m to k). Both forward and backward direction warps are used to address incorrect grouping due to occlusion. Because accurate warp estimation requires a segmentation, which is unknown, we use Classic-NL [10], robust to motion discontinuities, to roughly approximate the warps defined on Ω without a segmentation.

Object Detection: Algorithm 2 specifies Quickest Moving Object Detection. The algorithm re-estimates Λ_n online at each new arrival of I_n. At each new acquisition of I_n, the algorithm finds a

Algorithm 1 *Region segmentation and likelihood $\Lambda_{k,n}$ given change time k and time n.*

// *initialize R_m for gradient descent of E_{seg}*
Compute the Classic-NL warp $w_{i,i+1}^{NL} : \Omega \to \Omega$
Compute $w_{m,n}^{NL}, w_{m,k}^{NL}$ using (5) and cluster together into two regions, denoted $R_m^j, j = 0, 1$
repeat // *gradient descent of E_{seg} for region R_m*
    Compute R_i by propagating R_m frame by frame to frame i using [5]
    Solve for Sobolev warp, $w_{i,i+1}^j = \arg\min_w \int_{R_i^j} \text{Res}_i^j(x)\,\mathrm{d}x$, extended to Ω
    Compute $w_{m,i}^j$ for $i = k, \ldots, n$ using (5)
    Compute f^j and update R_m by a gradient step of E_{seg} in (13)
until R_m does not change between iterations
Compute Sobolev warps $w_{i,i+1}^{null} : \Omega \to \Omega$ for null hypothesis
Set $\text{Res}_i^{null}(x) = \rho(I_{i+1}(w_{i,i+1}^n(x)) - I_i(x))$ and $\text{Res}_i^j(x) = \rho(I_{i+1}(w_{i,i+1}^j(x)) - I_i(x))$
Compute $-\log \Lambda_{k,n} = \sum_{i=k}^n \int_\Omega \text{Res}_i^{null}(x)\,\mathrm{d}x - \sum_{j=0}^1 \int_{R_i^j} \text{Res}_i^j(x)\,\mathrm{d}x$
return $\Lambda_{k,n}$ and R_n as the proposed detection in frame n

change time $k \in \{2, \ldots, n-1\}$ by revisiting each previous data and solving an optimization problem (Lines 4-6), which is solved by Algorithm 1.

Algorithm 2 *Quickest Moving Object Detection and Segmentation*

Set $n = 1$
repeat // *compute likelihood ratio Λ_0*
    Increment $n \leftarrow n + 1$, acquire image I_n
    for $k = 2, \ldots, n-1$ **do** // *find a most probable change time k*
        Compute $\Lambda_{k,n}$ and $R_{k,n}$ using Algorithm 1
    end for
    Set $\Lambda_n = \Lambda_{k_n^*,n}$ and $R_n = R_{k_n^*,n}$ where $k_n^* = \arg\max_{2 \le k < n} \Lambda_{k,n}$
until $\Lambda_n > b$
return R_n as the object detection at time n

Simplification for Efficiency: Since the post-change distribution depends on hidden variables, data must be revisited as discussed in Section 2. Thus, in the general case, Lines 4-6 in Algorithm 2 cannot be simplified without possible loss of optimality. However, we propose a heuristic that estimates the change time k_n^* without having to explicitly evaluate $\Lambda_{k,n}$ for each k. This is done by applying QCD to simpler distributions than (9) and (10). If the residuals $\text{Res}_i^{NL}(x)$ are assumed iid and distributed according to $\mathcal{N}(\mu_0, \sigma_0)$ pre-change and $\mathcal{N}(\mu_1, \sigma_1)$ post-change, then the following statistic, which can be computed efficiently, arises

$$F_{k,n} = \sum_{i=k}^{n} \exp\left\{ -\frac{1}{|\Omega|} \int_\Omega \frac{(\text{Res}_i^{NL}(x) - \hat{\mu}_{1,k})^2}{\hat{\sigma}_{1,k}^2} - \frac{(\text{Res}_i^{NL}(x) - \hat{\mu}_{0,k})^2}{\hat{\sigma}_{0,k}^2}\,\mathrm{d}x \right\}, \tag{14}$$

where $\hat{\mu}_{0,k}, \hat{\sigma}_{0,k}$ are estimated from $\text{Res}_i^{NL}(x), i = 1, \ldots, k$ and $\hat{\mu}_{1,k}, \hat{\sigma}_{1,k}$ are estimated from $\text{Res}_i^{NL}(x), i = k+1, \ldots, n$. Therefore, we modify Algorithm 2 to Algorithm 3. Note that the k^* that achieves a maximum of $F_{k,n}$ is proposed as the maximizer of $\Lambda_{k,n}$. Only $\Lambda_{k^*,n}$, which approximates Λ_n and saves computation. Note that $F_{k^*,n}$ is not thresholded to obtain the stopping time. Although this would lead to additional computational savings, this is avoided since it is difficult to tune a threshold for a wide range of sequences. This is expected since the distributions that $F_{k,n}$ assumes are not as accurate as (9) and (10). However, the maximizer k^* is a useful candidate for the change time in estimating Λ_n.

4 Experiments

We now test our method in experiments and verify it against the theory of Quickest Change. Since we are not aware of an extensive dataset for the segmentation and detection of moving objects, which start to move or come into view at some unknown time in the video and continues moving, against

Algorithm 3 *Faster Quickest Moving Object Detection and Segmentation*

Set $n = 1$
repeat // *compute likelihood ratio* Λ_0
  Increment $n \leftarrow n + 1$, acquire image I_n
  for $k = 2, \ldots, n-1$ **do** // *find a most probable change time* k
    Compute $F_{k,n}$ as in (14)
  end for
  Set $k_n^* = \arg\max_{2 \le k < n} F_{k,n}$, compute $\Lambda_n = \Lambda_{k_n^*,n}$ and $R_n = R_{k_n^*,n}$ using Algorithm 1
until $\Lambda_n > b$
return R_n as the object detection at time n

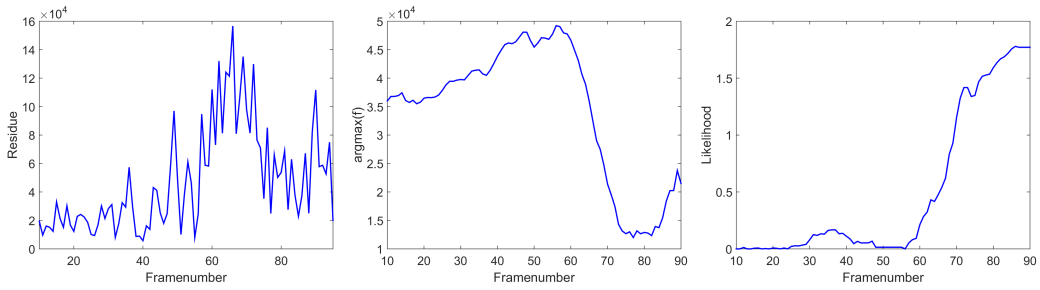

Figure 1: **Time evolution of statistics**. *We examine a image sequence where change happens at* $\Gamma = 58$. *[Left]: A plot of the mean value of the residual* Res_i^{NL} *from Classic-NL optical flow [10] versus frame number in the video, which is too noisy for detecting changes. [Middle]: A plot of the statistic* $F_{k_n^*,n}$ *versus the frame number* n. *This statistic is less noisy, but unreliable for thresholding across sequences. [Right]: The likelihood ratio* Λ_n *versus frame number* n, *which can be thresholded over a wide range to achieve correct detection.*

a possibly dynamic scene or moving camera, we collected our own. We have collected 20 videos, some that we collected from Youtube, and a number from the dataset [14], those that contain camera motion or dynamic backgrounds. Ground truth segmentations are available for the videos in [14], and we segmented the others. Each sequence contains 100-200 frames.

We are not aware of a competing online detection and segmentation approach from video that addresses detection with minimum delay, operating with false alarm constraints. According to the theory, the derived algorithm provides the optimal solution under the models assumed in Section 3.1. However, in practice due to the estimation of probabilities, which depend on hidden variables that are determined by local optimization methods, the algorithm may not be optimal. However, we show next that compared to a baseline approach of quickest change detection determined from thresholding the test statistic $F_{k,n}$ (14), our method achieves smaller delay for any false alarm constraint.

First, we illustrate the advantage of considering the likelihood ratio Λ_n over simpler statistics in Figure 1. There, the mean value of the frame-by-frame residual from Classic-NL is shown, which illustrates the difficulty in detecting the change due to the noise in the statistics. The plot of Λ_n versus time n shows a clean signal, and confirms our expected behavior of growing over after the change when more evidence is available. It can easily be thresholded over a wide range without false alarm. The simpler statistic $F_{k,n}$ is cleaner that the residual, but many times does not show the desired behavior, although its optimizer k_n^* (not shown in the plot), is useful in computing Λ_n.

Next we run our method over the whole dataset and different thresholds $b = 0.1, \ldots, 1.8$. Sample qualitative results of the segmentations and detections are shown in Figure 3. It generally shows fewer false alarms as the threshold increases, but with greater delay (when, for example, more of the object moves within the video). We quantity the results using a plot of the false alarm rate versus the average detection delay in Figure 2. The average detection delay is defined empirically as the average difference of the detection time and the ground truth change time, i.e., $n - \Gamma$ (if $n \ge \Gamma$) over all sequences. The empirical false alarm rate is the number of detections before the change time or segmentation (with less than 0.75 f-measure) after the change time divided by the number of sequences. The figure shows the trend of fewer false alarms at the cost of longer delay. Also, we show the results of a baseline change detection procedure using a thresholding of $F_{k,n}$, which has greater detection delay than our approach at any false alarm rate.

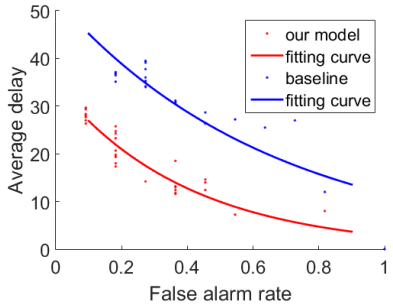

Figure 2: **Delay versus false alarm tradeoff**. *The average detection delay (over all sequences in our dataset) versus the false alarm rate as the threshold is varied. The blue curve uses thresholds of $F_{k_n^*,n}$ as a stopping time. The red curve uses our proposed stopping time from Λ_n. Our approach leads to less delay.*

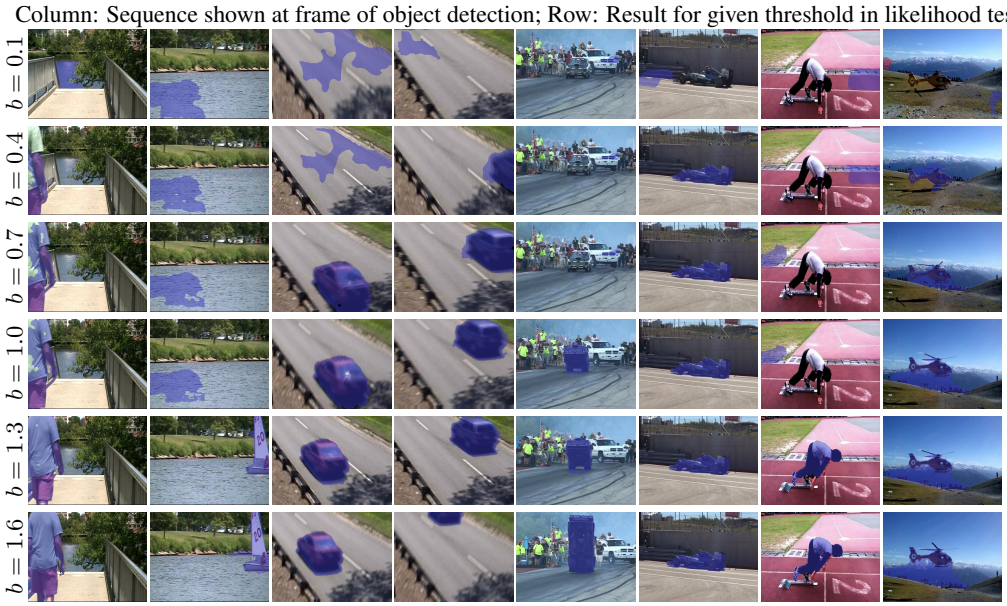
Column: Sequence shown at frame of object detection; Row: Result for given threshold in likelihood test
$b = 0.1$
$b = 0.4$
$b = 0.7$
$b = 1.0$
$b = 1.3$
$b = 1.6$

Figure 3: **Sample detections for varying thresholds of Λ_n**. *Each sequence contains one object (person, boat, car, or helicopter) moving at some unknown time with camera motion and dynamic background. Segmentations are purple masks. Small thresholds lead to several false alarms. As the threshold is increased, false alarms are mitigated, but with greater delay (number of frames after the object moves). The right column is a failure case, which occurs since the segmentation gets trapped in irrelevant background clutter.*

5 Conclusion

We have derived a method for moving object detection and segmentation in video under the assumption of dynamic backgrounds. The method was derived based on Quickest Change Detection, which led us to an online algorithm that minimizes the delay in detection subject to false alarm constraints. Although our problem leads to a non-recursive algorithm, where the data must be revisited at arrival of new data, we proposed a simplifying heuristic that led to cost savings and was used effectively in experiments. The experiments showed that our algorithm achieved less delay under various false alarm constraints compared to a baseline approach. Future work will look into further speed-ups.

References

[1] P. Ochs, J. Malik, and T. Brox, "Segmentation of moving objects by long term video analysis," *PAMI*, vol. 36, no. 6, pp. 1187–1200, 2014. 1, 2

[2] J. Y. Wang and E. H. Adelson, "Representing moving images with layers," *IEEE TIP*, vol. 3, no. 5, pp. 625–638, 1994. 1, 2

[3] Y. Weiss and E. H. Adelson, "A unified mixture framework for motion segmentation: Incorporating spatial coherence and estimating the number of models," in *Computer Vision and Pattern Recognition, 1996. Proceedings CVPR'96, 1996 IEEE Computer Society Conference on*. IEEE, 1996, pp. 321–326. 1, 2

[4] D. Sun, J. Wulff, E. B. Sudderth, H. Pfister, and M. J. Black, "A fully-connected layered model of foreground and background flow," in *CVPR*. IEEE, 2013, pp. 2451–2458. 1, 2

[5] Y. Yang, G. Sundaramoorthi, and S. Soatto, "Self-occlusions and disocclusions in causal video object segmentation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 4408–4416. 1, 2, 5, 6

[6] H. V. Poor and O. Hadjiliadis, *Quickest detection*. Cambridge University Press Cambridge, 2009, vol. 40. 2

[7] V. V. Veeravalli and T. Banerjee, "Quickest change detection," *Academic press library in signal processing: Array and statistical signal processing*, vol. 3, pp. 209–256, 2013. 2, 3

[8] B. D. Lucas, T. Kanade *et al.*, "An iterative image registration technique with an application to stereo vision." in *IJCAI*, vol. 81, 1981, pp. 674–679. 2

[9] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert, "High accuracy optical flow estimation based on a theory for warping," in *ECCV*. Springer, 2004, pp. 25–36. 2

[10] D. Sun, S. Roth, and M. J. Black, "Secrets of optical flow estimation and their principles," in *CVPR*. IEEE, 2010, pp. 2432–2439. 2, 4, 5, 7

[11] A. Ayvaci, M. Raptis, and S. Soatto, "Occlusion detection and motion estimation with convex optimization," in *Advances in neural information processing systems*, 2010, pp. 100–108. 2

[12] D. Sun, E. B. Sudderth, and M. J. Black, "Layered image motion with explicit occlusions, temporal consistency, and depth ordering," in *Advances in Neural Information Processing Systems*, 2010, pp. 2226–2234. 2

[13] T. B. M. Keuper, B. Andres, "Motion trajectory segmentation via minimum cost multicuts," *IEEE International Conference on Computer Vision (ICCV)*, pp. 3271–3279, 2015. 2

[14] N. Goyette, P.-M. Jodoin, F. Porikli, J. Konrad, and P. Ishwar, "A novel video dataset for change detection benchmarking," *Image Processing, IEEE Transactions on*, vol. 23, no. 11, pp. 4663–4679, 2014. 2, 7

[15] Y. Benezeth, P.-M. Jodoin, B. Emile, H. Laurent, and C. Rosenberger, "Comparative study of background subtraction algorithms," *Journal of Electronic Imaging*, vol. 19, no. 3, pp. 033 003–033 003, 2010. 2

[16] I. Cohen and G. Medioni, "Detecting and tracking moving objects for video surveillance," in *Computer Vision and Pattern Recognition, 1999. IEEE Computer Society Conference on.*, vol. 2. IEEE, 1999. 2